Inferring Autonomous System Relationships in the Internet

Lixin Gao
Dept. of Electrical and Computer Engineering
University of Massachusetts, Amherst
http://www-unix.ecs.umass.edu/~lgao

Outline

- Internet Architecture and Routing
- AS Relationships
- Heuristic Algorithms
- Experimental Results

AS Commercial Relationships

- Provider-customer:
  - customer pays its provider for transit services
- Peer-peer:
  - exchange traffic between customers
  - no money exchange
- Sibling-sibling:
  - have mutual transit agreement
  - merging ISPs, Internet connection backup

Route Propagation Policy

- Constrained by contractual commercial agreements between administrative domains

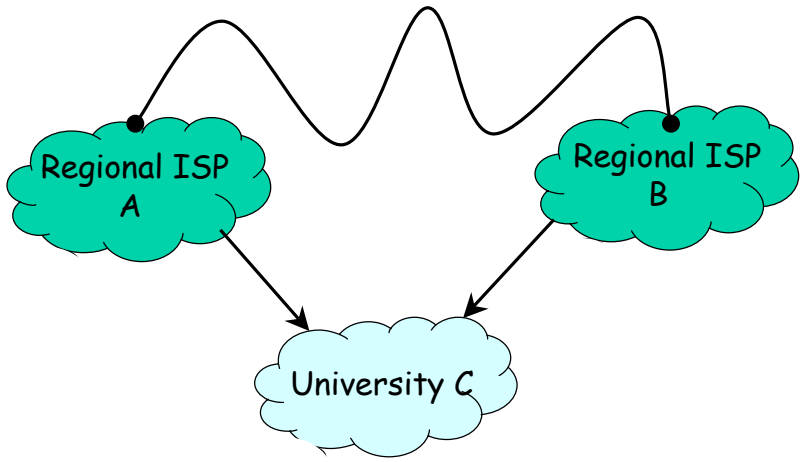

e.g., An AS does not provide transit services between its providers

Why Infer AS Relationships?

- Crucial part of Internet structure
  - Connectivity does not imply reachability
  - Connectivity alone can not fully characterize structural properties of Internet
- No registry of AS relationships
  - Many ISPs are not willing to reveal their relationships to others in IRR
  - Relationships are evolving; hard to be up-to-date

Applications of AS Relationships

- Construct distance map
- Place proxy or mirror site servers
- Potentially avoid route divergence

- Help ISPs or domain administrators to achieve load balancing and congestion avoidance
- Help ISPs or companies to plan for future contractual agreements
- Help ISPs to reduce effect of misconfiguration and to debug router configuration files

AS Relationship Graph

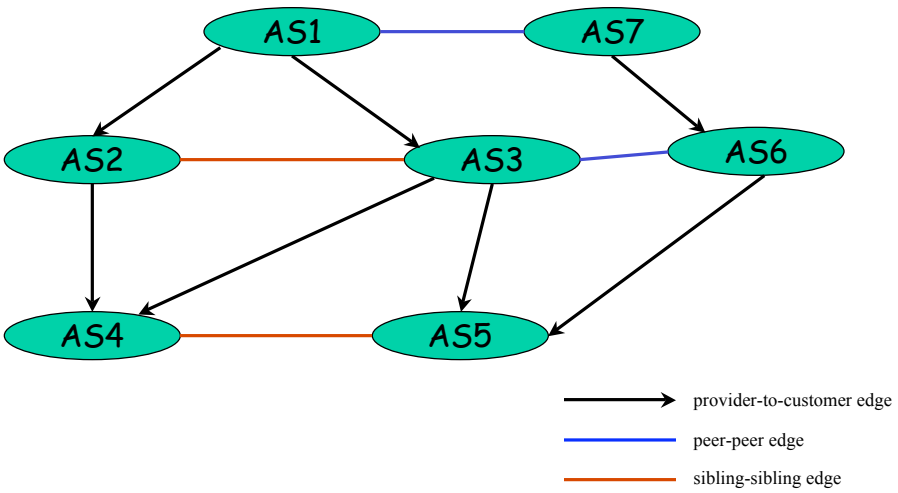

Route Propagation Rule

- An AS or a set of ASes with sibling relationship does not provide transit services between any two of its providers and peers
- BGP routing table entries have certain patterns

Routing Table Entry

Network	Next hop	AS Path
4.2.24.0/21	134.24.127.3	1740 1 i
	194.68.130.254	5459 5413 1 i
	158.43.133.48	1849 704 702 701 1 i
	193.0.0.242	3333 286 1 i
	144.228.240.93	1239 1 i

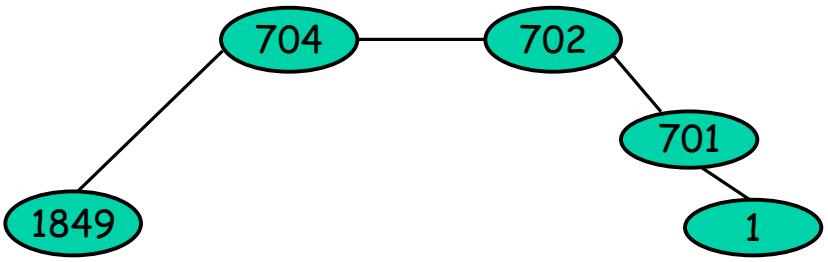

Routing Table Entry Patterns

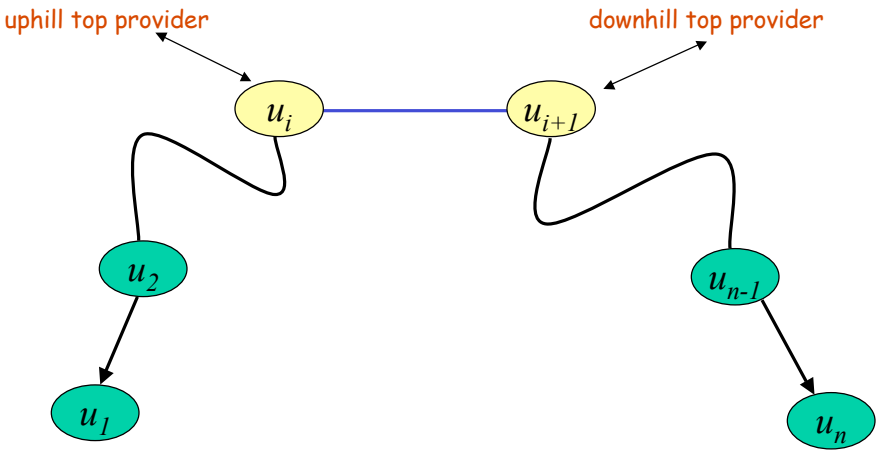

Heuristic Algorithms

- Infer provider-customer and sibling-sibling
  - basic
  - refined
- Infer peer-peer
  - final

Basic Algorithms

- Heuristics:
  - Top provider has largest degree
- Based on patterns on BGP routing table entries
  - Consecutive AS pairs on the left of top provider are customer-to-provider or sibling-sibling edges
  - Consecutive AS pairs on the right of top provider are provider-to-customer or sibling-sibling edges

Initialize Consecutive AS Pair Relationship

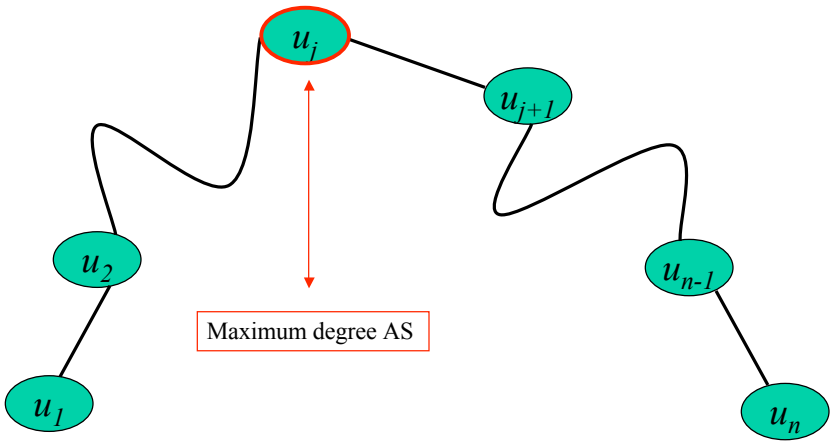

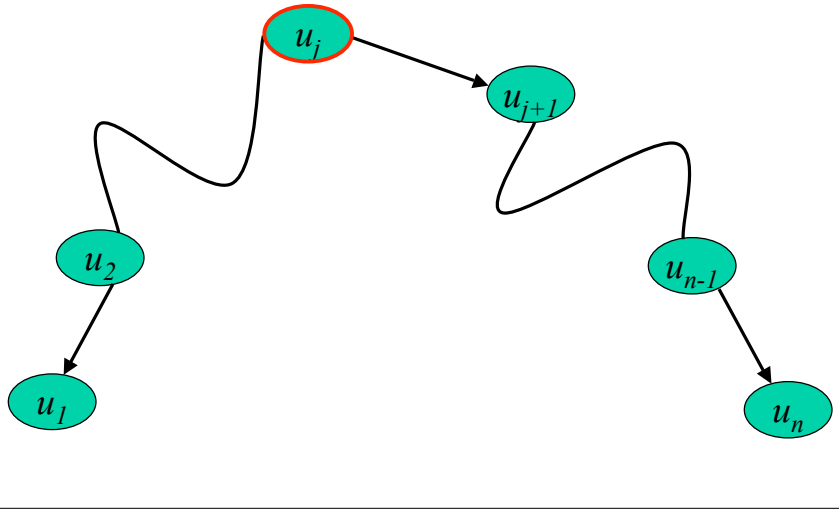

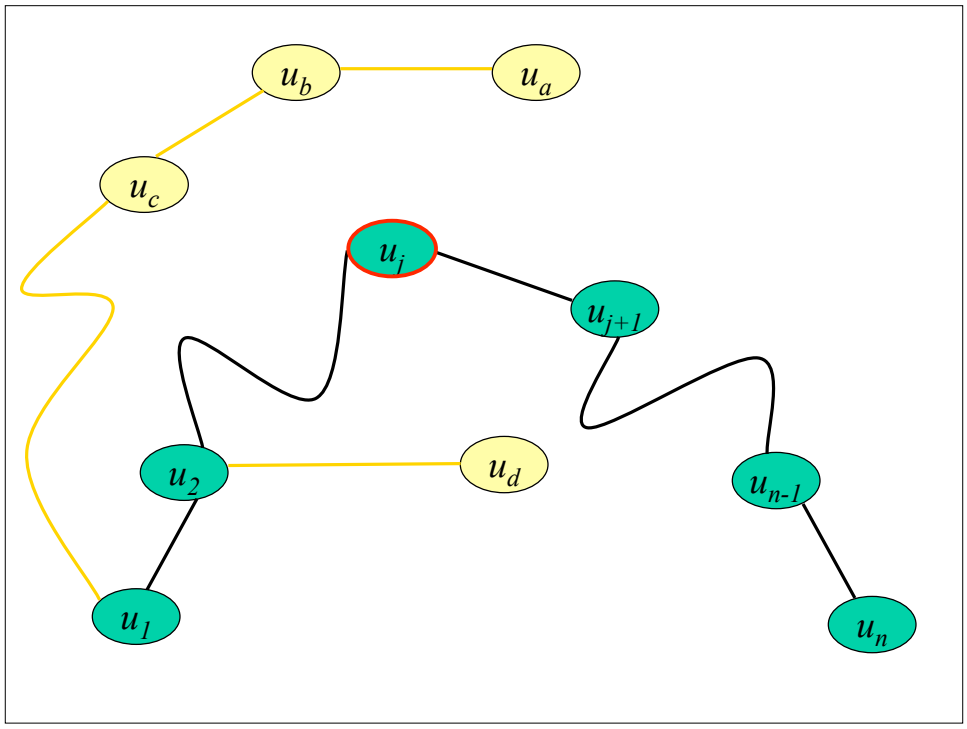
u_b
u_a
u_c
u_j
u_{j+1}
u_2
u_d
u_{n-1}
u_1
u_n

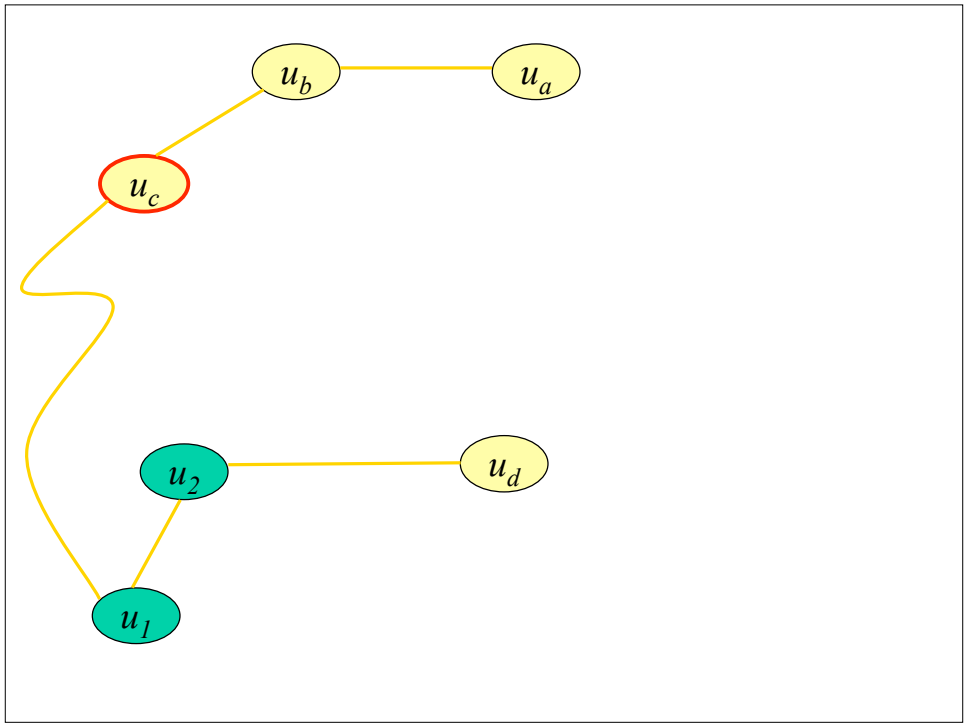
u_b
u_a
u_c
u_2
u_d
u_1

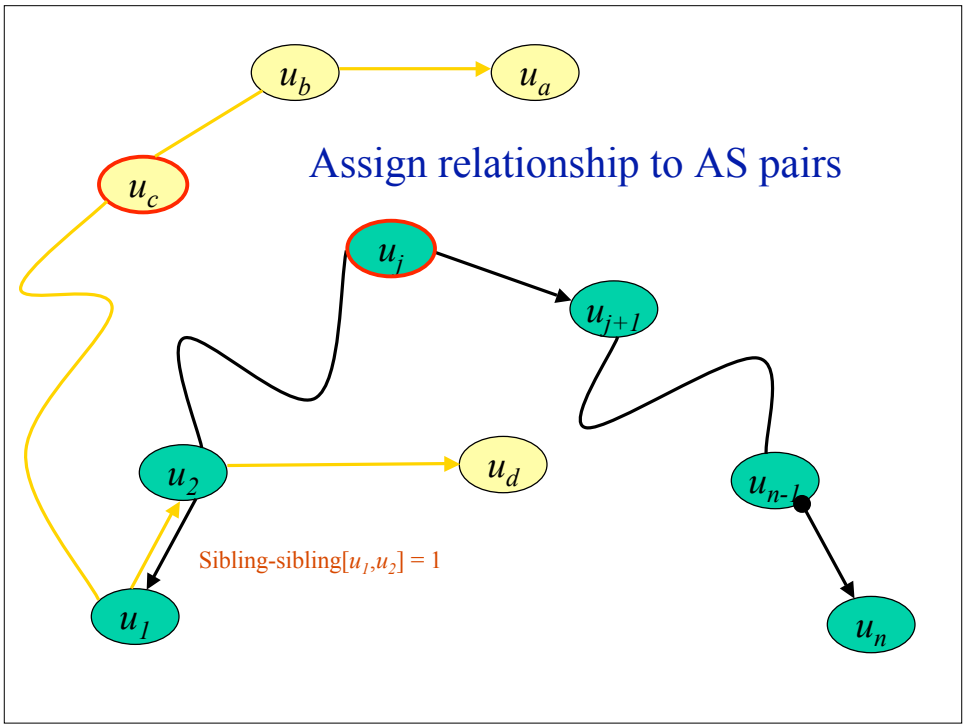

Refined Algorithm

- Bogus Routing Entries
- Each routing table entry votes on AS relationships
- Ignore sibling-to-sibling relationship concluded by only one entry

Inferring Peer-Peer Relationships

- Peer-peer edge is between top provider and one of its neighbors only
- Heuristics:
  - peer-to-peer edge is between top provider and its higher degree neighbor
  - degrees of two peers do not differ significantly
    - < **R** times

Final Algorithm

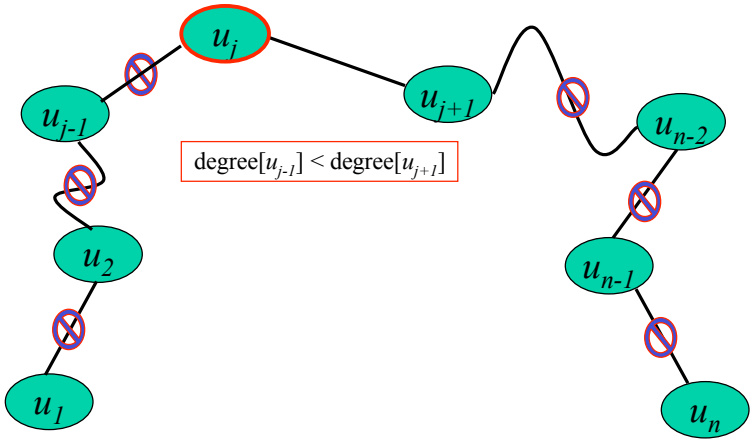

Final Algorithm

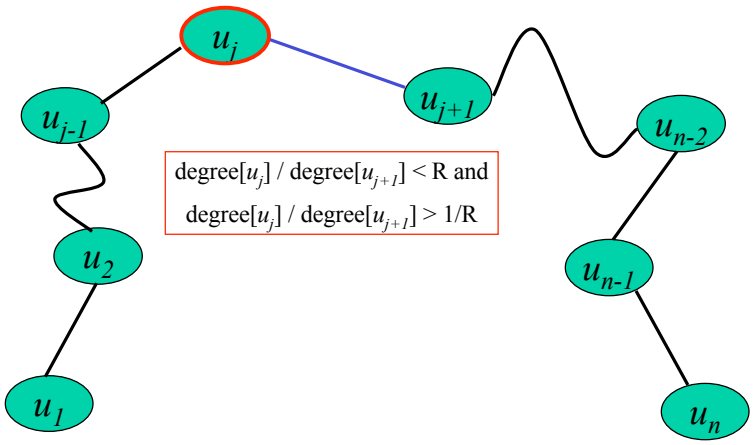

Experimental Verification

- Routing table from Route Views
  - Connected to 22 ISPs at 24 locations
  - Daily routing table dump
- Routing table from 3 days
  - 1999/9/27, 2000/1/2, 2000/3/9
  - ~1 million routing entries

Inference Results

	TOTAL ROUTING ENTRIES	TOTAL EDGES	SIBLING - SIBLING EDGES INFERRED BY BASIC (PERCENTAGE)	SIBLING - SIBLING EDGES INFERRED BY REFINED (IGNORED ENTRIES)	PEER-PEER EDGES INFERRED BY FINAL [R=INFINITY] (PERCENTAGE)	PEER-PEER EDGES INFERRED BY FINAL [R=60] (PERCENTAGE)
1999/9/27	968674	11288	149 (1.3%)	124 (25)	884 (7.8%)	733 (6.5%)
2000/1/2	936058	12571	186 (1.47%)	135 (51)	838 (6.7%)	668 (5.3%)
2000/3/9	1227596	13800	203 (1.47%)	157 (46)	857 (6.2%)	713 (5.7%)

Verification of Inferred Relationships by AT&T

OUR INFERENCE	AT&T INFORMATION	PERCENTAGE OF AS
Customer	Customer	99.8%
	Peer	0.2%
Peer	Peer	76.5%
	Customer	23.5%
Sibling	Sibling	20%
	Peer	60%
	Customer	20%
Nonexistent	Customer	95.6%
	Peer	4.4%

Comparing inference results from Basic and Final($R=\infty$) with AT&T internal information

Verification of Inferred Relationships by AT&T

OUR INFERENCE	AT&T INFORMATION	PERCENTAGE OF AS
Customer	Customer	99.5%
	Peer	0.5%
Peer	Peer	76.5%
	Customer	23.5%
Sibling	Sibling	25%
	Peer	50%
	Customer	25%
Nonexistent	Customer	95.6%
	Peer	4.4%

Comparing inference results from Refined and Final($R=\infty$) with AT&T internal information

Verification of Inferred Relationships by AT&T

OUR INFERENCE	AT&T INFORMATION	PERCENTAGE OF AS
Customer	Customer	99.8%
	Peer	0.2%
Peer	Peer	100%
Sibling	Sibling	20%
	Peer	60%
	Customer	20%
Nonexistent	Customer	95.6%
	Peer	4.4%

Comparing inference results from Basic and Final($R=60$) with AT&T internal information

WHOIS Lookup Service

- Supplies name and address of company that owns an AS
- AS pair might have sibling-sibling relation if
  - belong to the same company or two merging companies
  - belong to two small companies located closeby

Verification by WHOIS lookup Service

- Confirm 101 of 186 inferred sibling-sibling relationships (> 50%)
- Some unconfirmed sibling-sibling relationships might be due to
  - WHOIS service is not up to date
  - Not enough information
- Bogus Routes:
  - Router configuration typo: 7018 3561 7057 7075 7057
  - Misconfiguration of small ISPs:1239 11116 701 7018
  - ...

Conclusions and Further Work

- AS relationships are inherent aspect of Internet architecture
- Our heuristic algorithm is based on routing entry pattern derived from policy rules
- Verification:
  - AT&T (99%)
  - Whois services (>50%)
- Further Work:
  - Policy effect on AS path length
  - AS relationship evolution